

TECHNIQUE FOR ETHERNET ACCESS TO PACKET-BASED SERVICES

Technical Field

This invention relates to a technique enabling access to packet-based services, such as IP, Frame Relay, and ATM, through an Ethernet Protocol network.

Background Art

Presently, communication service providers, such as AT&T, offer high-speed data communications service to customers through a variety of access mechanisms. For example, a customer may gain network access through a private line connection, i.e., a direct link to the communications service provider's network. Private line access provides a dedicated port not shared by anyone else with facility bandwidth available exclusively to the particular customer. Unfortunately, private line access is expensive, and is practical only for customers that have very high traffic capacity requirements.

As an alternative to private line access, communications service providers such as AT&T also offer virtual circuit access allowing several customers to logically share a single circuit, thus reducing costs. Such shared circuits, typically referred to as Permanent Virtual Circuits, allow communication service providers to guarantee customer traffic flows that are distinguishable from each, secure, and allow customers to enjoy different service features. An example of such a technique for offering such shared service in a Multi-Protocol Label Switching Network is disclosed in U.S. Patent 6,081,524, assigned to AT&T.

Presently, there is a trend towards using Ethernet networks in place of Frame Relay and ATM networks especially for transporting traffic among two or more premises belonging to the same customer. Ethernet-based Metropolitan Area Networks (MANs) currently exist in many areas and offer significant cost advantages on a per port basis, as compared to Frame Relay and ATM service. Transmission speeds as high as 100, 1000 or even 10,000 MB/second are possible with such Ethernet MANs. Moreover, optical Ethernet MANs typically offer a rich set of features, flexible topology and simple-end-to-end provisioning.

Present-day Ethernet-based MANs lack the ability to logically separate traffic received from different customers, giving rise to issues of data security. Moreover, such present day Ethernet-based MANs lack the ability to manage bandwidth among customers, thus preventing the network from regulating customer traffic to assure equitable access. Thus, there is a need for a technique for routing data in an Ethernet protocol network that overcomes the aforementioned disadvantages.

Brief Summary of the Invention

Briefly, in accordance with a preferred embodiment, a method is provided for routing data in an Ethernet protocol network having a plurality of platforms, each serving one or more customers. A first platform receives at least one frame from a sending site (e.g., a first customer's premises) that is destined for a receiving site (e.g., another premises belonging to the same or a different customer.) After receiving the frame, the first platform overwrites a portion of the frame with a **customer descriptor** that specifically identifies the sending customer. In practice, the first platform will overwrite a Virtual Local Area Network (VLAN) field that is typically employed by the sending customer for the purposes of routing data among various VLANs at the sending premises. Rather than overwrite the VLAN field in the frame, the first platform could overwrite another field, such as the source address field, or could even employ a "shim" header containing the customer descriptor. All further use of the term **customer descriptor** implies that any of the above or similar techniques could be used.

After overwriting the frame with the customer descriptor, the sending platform routes the frame onto the MAN network for routing among the other platforms, thereby sharing trunk bandwidth and other resources, but logically distinct from other customer's traffic with different customer descriptors. A destination address in the frame directs the frame to its corresponding receiving platform. Upon receipt of the frame, the receiving platform uses the customer descriptor to segregate the frame for delivery to the proper destination, especially in the event where different customers served by the same platform use overlapping addressing plans. Thus, the customer descriptor in each frame

advantageously enables the receiving platform to distinguish between different customers served by that platform.

For traffic with a destination beyond the MAN, this method provides a convenient and efficient way to map the customer-descriptor to similar identifiers in a Wide Area Network (WAN), such as a Permanent Virtual Circuit (PVC), a Virtual Private Network (VPN), or an MPLS Label Switched Circuit.

Overwriting each frame with the customer-descriptor thus affords the ability to logically segregate traffic on the Ethernet MAN to provide Virtual Private Network (VPN) services of the type offered only on more expensive Frame Relay and ATM networks. Moreover, the customer descriptor used to tag each frame can advantageously include Quality of Service (QoS) information, allowing the sender to specify different QoS levels for different traffic types, based on the Service Level Agreement (SLA) between the customer and the communications service provider.

Brief Description of the Drawings

FIGURE 1 depicts an Ethernet Protocol Metropolitan Area Network (MAN) in which each information frame is tagged with a customer descriptor in its VLAN field in accordance with the present principles;

FIGURE 2 illustrates a sample information frame for transmission over the network of FIG. 1;

FIGURE 3 illustrates a portion of the MAN showing the various stages in the tagging process;

FIGURE 4 illustrates a portion of a MAN showing the use of the priority bits within the VLAN field to establish different Quality of Service levels;

FIGURE 5 illustrates a portion of a MAN showing the manner in which information frames are mapped to different Permanent Virtual Circuits by an ATM switch;

FIGURE 6 illustrates a portion of a MAN showing the manner in which information frames are mapped into different Multi-Protocol Label Switching (MPLS) tunnels; and

FIGURE 7 illustrates a portion of a MAN showing the manner in which information frames are mapped into different service networks.

Detailed Description

FIGURE 1 depicts an Ethernet Protocol Metropolitan Area Network (MAN) 10 comprised of a plurality of Multi-Service Platforms (MSPs) 12_1-12_n where n is an integer, each MSP taking the form of an Ethernet switch or the like. In the illustrated embodiment $n=4$ although the network 10 could include a smaller or larger number of MSPs. A fiber ring or SONET ring infrastructure 14 connects the platforms 12_1-12_4 in daisy-chain fashion allowing each MSP to statistically multiplex information onto, and to statistically de-multiplexing information off the ring infrastructure 14.

Each of MSPs 12_1-12_3 serves at least one, and in some instances, a plurality of premises 16 belonging to one or more customers. In the illustrated embodiment of FIG. 1, the MSP 12_1 serves a single customer premises 16_1 belonging to customer 1 whereas, the MSP 12_2 serves premises 16_2 , and 16_3 belonging to customers 2 and 3, respectively. The MSP 12_3 serves a single premises 16_4 that belongs to customer 3. The MSPs 12_1-12_3 are linked to their corresponding premises via 10, 100 or 1000 MB links 18. The MSP 12_4 bears the legend "CO MSP" because it serves as a central office to route traffic from the MAN 10 to a Provider Edge Router (PER) 18 for delivery to other networks, such as Frame Relay, ATM, MPLS networks or the Internet as discussed hereinafter. By the same token, the PER 18 can route traffic from such other networks onto the MAN 10 via the CO MSP 12_4 .

The traffic routed onto and off of the MAN 10 by each MSP takes the form of one or more information frames 20 depicted in FIG. 2. Heretofore, traffic routed onto the network 10 from a particular customer's premises was combined with other customer's traffic with no logical separation, thus raising security concerns. Moreover, since all customers' traffic share the same bandwidth, difficulties existed in prior art Ethernet MANs in regulating the traffic from each customer's premises, and in affording different customers different Quality of Service (QoS) level in accordance with individual Service Level Agreements.

These difficulties are overcome in accordance with the present principles by “tagging” each frame 20 routed onto the network 10 at a particular MSP, say MSP 12₃, with a customer descriptor 22’ (best seen in FIG. 2) that identifies the customer sending that frame. As discussed in greater detail below, each MSP receiving a frame 20 on the fiber ring infrastructure 14 uses the customer descriptor 22’ in that frame to maintain distinct routing and addressing tables that are assigned to each customer served by that MSP. This permits each customer to use their own addressing plan without fear of overlap with other customers, as they are all maintained as logically separate.

FIGURE 2 depicts the structure of an exemplary Ethernet protocol frame 20 specified by Ethernet Standard 802.1Q. Among the blocks of bytes within each frame 20 is a Virtual Local Area Network (VLAN) Identifier 22 comprised of sixteen bits. In practice, the VLAN identifier 22, in conjunction with a VLAN flag block 23 within the frame, facilitates routing of the frame within a customer’s premises to a particular VLAN. However, the VLAN identifier 22 has no influence on routing of the frame 20 after receipt at a MSP.

In accordance with the present principles, the prior disadvantages associated with conventional Ethernet networks, namely the lack of security and inability to regulate QoS levels, are overcome by overwriting the VLAN identifier 22 in each frame 20 with the customer descriptor maintained by the service provider. Overwriting the VLAN identifier 22 of FIG. 2 of each frame 20 with the customer descriptor 22’ serves to “tag” that frame with identity of its sending customer identity, thus affording each MSP in the network 10 the ability to route those frames only among the premises belonging to that same sending customer. Such tagging affords the operator of the network 10 the ability to provide security in connection with frames transmitted across the network, since frames with customer ID A would not be delivered to any premises of customer with ID B. As an example, two or more customers served by a single MSP may use overlapping IP addressing schemes. In the absence of any other identifier, the MSP receiving such frame lacks the ability to assure accurate delivery.

In the illustrated embodiment depicted in FIG. 2, each MSP of Fig. 1 tags the frame 20 by overwriting the VLAN identifier 22 with the customer descriptor 22’. However, tagging could occur in other ways, rather than overwriting the VLAN identifier

22. For example, the source address block 25 within the frame 20 could be overwritten with the customer descriptor 22'. Alternatively, the data field 25 could include a shim header comprising the customer descriptor 22'.

The tagging of each frame 20 with the customer descriptor 22' affords several distinct advantages in connection with routing of the frames through the MAN. First, as discussed above, the tagging affords each recipient MSP the ability to distinguish traffic destined for customers with overlapping address schemes, and thus allows for segregating traffic on the MAN. Further, tagging each frame 20 with the customer descriptor 22' enables mapping of the frames from a MAN 100 depicted in FIG 3 to corresponding one of a plurality of customer Virtual Private Networks 26₁-26₃ within an MPLS network 28. As seen in FIG. 3, an MSP 120₂ within the MAN 100 receives traffic from premises 160₁, 160₂, and 160₃ belonging to customer 1, customer 2 and customer 3, respectively, which enjoy separate physical links to the MSP. Upon receipt of each from a particular customer, the MSP 120₂ overwrites that frame with the customer descriptor corresponding to the sending customer.

After tagging each data frame, the MSP 120₂ statistically multiplexes the frames onto the fiber ring infrastructure 14 for transmission to a CO MSP 120₄ for receipt at a destination PER 180 that serves the MPLS network 28 within which are customer Virtual Private Networks 26₁-26₃. Using the customer descriptor in each frame, the PER 180 maps the frame to the corresponding VPN identifier associated with a particular one of customer Virtual Private Networks 26₁-26₃ to properly route each frame to its intended destination.

The tagging scheme of the present invention also affords the ability to route information frames with different QoS levels within a MAN 1000 depicted in FIG 4. As seen in FIG. 4, an MSP 1200₂ within the MAN 1000 receives traffic from premises 1600₂, and 1600₃ belonging to customer 2 and customer 3, respectively, which enjoy separate physical links to the MSP, allowing each to send information frames into the MAN. In the illustrated embodiment of FIG. 4, the frames originating from the premise 1600₂ may contain either voice or data and have corresponding QoS level associated with each type of frame. Upon receiving such frames, the MSP 1200₂ overwrites the frame with the customer descriptor corresponding to the particular customer sending the frame.

The customer descriptor will not only contain the identity of the sending customer, but the corresponding QoS level associated with that frame.

After tagging each data frame, the MSP 1200₂ statistically multiplexes the frames onto the fiber ring infrastructure 14 for transmission to a CO MSP 1200₄ for receipt at a destination PER 1800 that serves an MPLS network 280 within which are customer Virtual Private Networks 260₂ and 260₃. Using the customer descriptor in each frame, the PER 1800 of FIG. 4 maps the frame to the corresponding customer VPN to properly route each frame to its intended customer premises. Further, the PER 1800 of FIG. 4 also maps the QoS level specified in the customer descriptor in the frame to assure that the appropriate quality of service level is applied to the particular frame.

In the above-described embodiments, the frames of customer traffic have been assumed to comprise IP packets that terminate on a router (i.e., Provider Edge Routers 18, 180 and 1800) and that the VPNs employ MPLS-BGP protocols. However, some customers may require multi-protocol support, or may otherwise require conventional PVCs so that the traffic streams must be mapped into Frame Relay or ATM PVCs as depicted in FIG. 5, which illustrates a portion of a MAN 10000 having a CO MSP 12000₄ serving an ATM switch 30 that receives traffic from the MAN. As seen in FIG. 5, each of premises 16000₁, 16000₂ and 16000₃ belonging to customer 1, customer 2 and customer 3, respectively 5 may originate information frames for receipt at MSP 12000₂ in the MAN 10000. The MSP 12000₂ tags each frame with the corresponding customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 to the CO MSP 12000₄ for receipt at the ATM switch 30. The ATM switch 30 then maps the frame to the appropriate PVC in accordance with the customer descriptor in the frame in a manner similar to the mapping described with respect to FIG. 3. Thus, the ATM switch 30 could map the frame to one of Frame Relay recipients' 32₁, 32₂, or 32₃, ATM recipients 32₄ or 32₅ or IMA (Inverse Multiplexing over ATM) recipient 32₆.

FIG. 6 depicts a portion of a MAN network 100000 that routes frames onto separate MPLS tunnels 40₁-40₃ (each emulating a private line 32 in an MPLS network 2800) in accordance with the customer descriptor written into each frame by a MSP 120000₂ in the MAN. Each of customer premises 160000₁, 160000₂ and 160000₃

depicted in FIG. 6 originate information frames for receipt at MSP 120000₂. The MSP 120000₂ tags each frame with the customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 for delivery to a CO MSP 120000₄ that serves a PER 18000. The PER 18000 translates the customer descriptors written onto the frames by the MSP 120000₂ into the MPLS tunnels 40₁-40₃ to enable the PER to route the traffic to the intended customer.

FIG. 7 depicts a portion of a MAN network 1000000 for routing traffic (i.e., information frames) onto separate networks in accordance with the customer descriptor written into each the frame by a MSP 120000₂ in the MAN. Each of customer premises 1600000₂ and 1600000₃ depicted in FIG. 7 originates information frames for receipt by the MSP 1200000₂. The MSP 1200000₂ tags each frame with the customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 for delivery to a CO MSP 1200000₄ that serves a PER 180000. In accordance with the customer descriptor, the PER 1800000 of FIG. 7 routes traffic to a particular one of several different networks, e.g., an Intranet VPN 42₁, a voice network 42₂ and the Internet 42₃, in accordance with the customer descriptor written onto the frame by the MSP 12₂.

The above-described embodiments merely illustrate the principles of the invention. Those skilled in the art may make various modifications and changes that will embody the principles of the invention and fall within the spirit and scope thereof.